1992

Lawrence Livermore
National Laboratory
A Heritage of Excellence
A COMMITMENT
TO THE FUTURE

NERSC

National Energy Research Supercomputer Center Machine Room

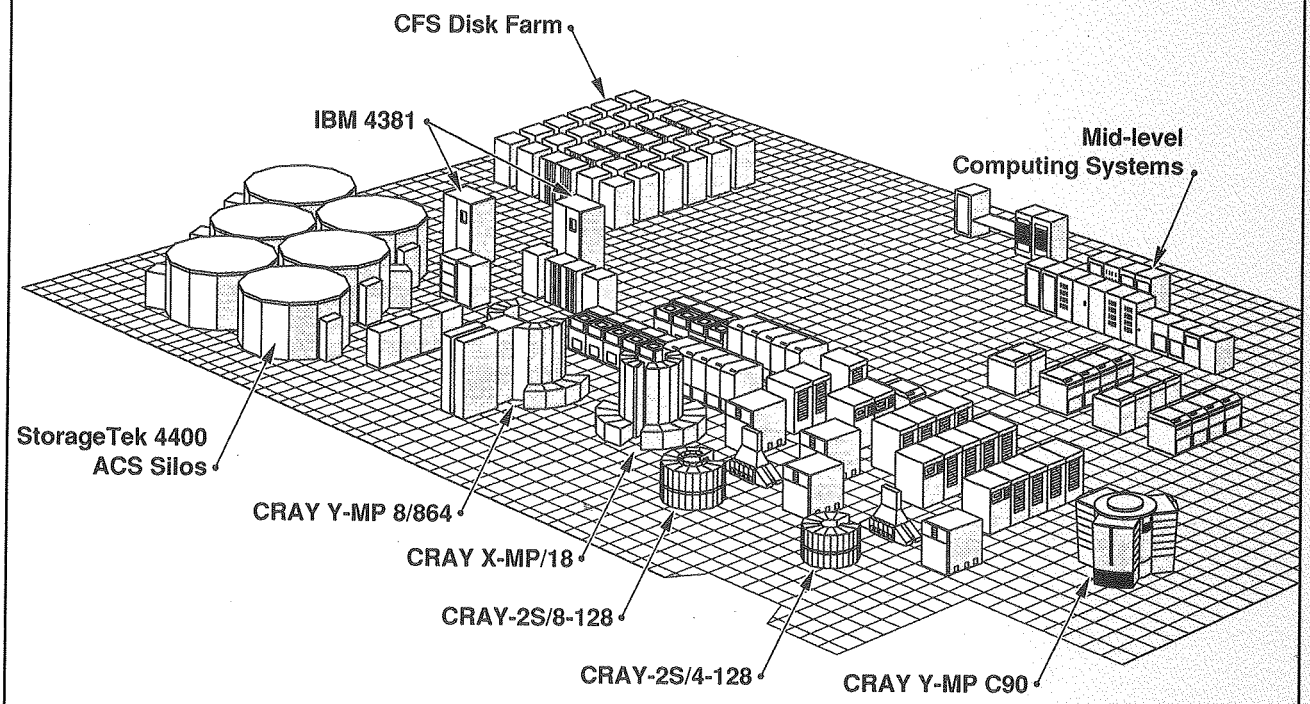

Building 451 • Lawrence Livermore National Laboratory • Livermore CA

National Energy Research Supercomputer Center Machine Room

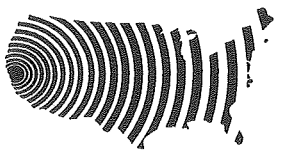

The National Energy Research Supercomputer Center (NERSC) was established for unclassified research in 1974. The Center provides large-scale computational support to the Energy Research community of the U.S. Department of Energy (DOE). NERSC applies advanced computer technology and computing techniques to problems in the areas of magnetic fusion energy, basic energy science, high energy and nuclear physics, health and environmental sciences, applied mathematical sciences, and Superconducting Super Collider research.

NERSC provides supercomputing services to about 4500 scientists and researchers at nearly 150 institutions throughout the country, including 26 government research laboratories, 92 universities, 11 private laboratories, and 19 industrial sites.

The Center also makes its services available to scientists in other countries. Committed to the belief that pooling resources furthers progress, NERSC encourages communication among American energy researchers and with researchers in other countries. The Center facilitates the sharing of information, codes, data, manpower, and computer power.

NERSC's computing power is concentrated in the Machine Room. This guide to the Machine Room describes four categories of computing facilities:

Large Computer Systems
Data Storage Systems
Mid-Level Computing Systems
Communications Systems

Most of these facilities are shown in the drawing on the cover. This guide also describes other important support services: Machine-Room operations, archival records, fire protection, environmental and safety monitoring; and prime power, cooling, and heating.

Large Computer Systems

Our large computer systems are composed of several CRAY supercomputers. These computers are the biggest and fastest in use today. The CRAY-2s, CRAY Y-MP, and CRAY Y-MP C90 perform scientific calculations for energy researchers throughout the United States and other countries. Our CRAY X-MP is the National Education Supercomputer (NES) and is dedicated to educational use by students and teachers.

CRAY-2 Systems (2)

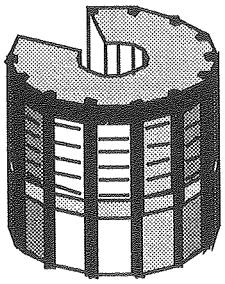

	CRAY-2S/4-128 Serial No. 2018	CRAY-2S/8-128 Serial No. 2101
No. CPUs (central processing units)	4	8 (only CRAY-2 with 8)
CPU theoretical performance	500 million arithmetic operations per second per CPU	500 million arithmetic operations per second per CPU
Memory capacity	>1 billion characters	>1 billion characters
Disk storage capacity	101 billion characters	84 billion characters
Disk storage cost	$119/1 million chars	$119/1 million chars
Power consumption	180 kilowatts electricity	180 kilowatts electricity
Cycle time	4.1 nanoseconds	4.1 nanoseconds
Coolant	250 gal circulating fluorinert:	250 gal circulating fluorinert:
Computer Reservoir, pipes Cost	150 gal 100 gal $200/gal	150 gal 100 gal $200/gal
Manufacturer	Cray Research, Inc.	Cray Research, Inc.
Installation date	August 1988	April 1990 (last one manufactured)
Machine cost	$17.5 million	$19 million

CRAY X-MP/18 Serial No. 312

This is the National Education Supercomputer, used by high school students and teachers in the DOE's National Education Supercomputer Program.

No. CPUs	1
Performance	210 million arithmetic operations per second
Memory capacity	64 million characters
Disk storage capacity	>4 billion characters
Disk storage cost	$212 per 1 million characters
Cycle time	9.5 nanoseconds
Manufacturer	Cray Research, Inc.
Installation date	October 1990
Machine cost	Donated by Cray Research, Inc.

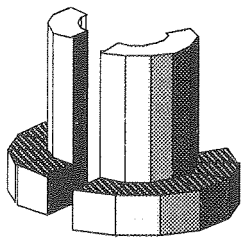

CRAY Y-MP 8D/864 Serial No. 1003

No. CPUs	8
Theoretical peak speed	333 million arithmetic operations per second per CPU
Memory capacity	512 million characters
Disk storage capacity	74 billion characters
Disk storage cost	$119 per 1 million characters
1 solid-state disk (SSD)	>2 billion characters capacity
Power consumption	245 kilowatts of electricity
Cycle time	6.3 nanoseconds
Manufacturer	Cray Research, Inc.
Installation date	March 1992
Machine cost	$18 million

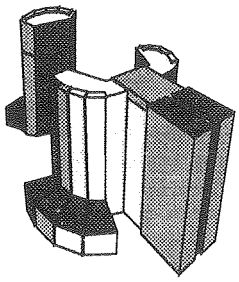

CRAY Y-MP C90/16256 Serial No. 4005

No. CPUs	16
Theoretical peak speed	1 billion arithmetic operations per second per CPU
Memory capacity	>2 billion characters
Disk storage capacity	113 billion characters
Disk storage cost	$119 per 1 million characters
Power consumption	368 kilowatts of electricity
Cycle time	4 nanoseconds
1 solid-state disk (SSD)	>4 billion characters capacity
Manufacturer	Cray Research, Inc.
Installation date	September 1992
Machine cost	$30 million

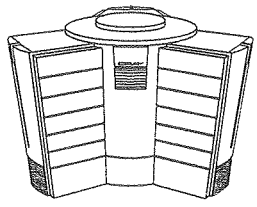

Data Storage Systems

Common File System (CFS)

CFS processes all our users' requests for data storage and retrieval. There are two IBM 4381s with their associated disks, and six StorageTek Automated Cartridge Systems serving CFS.

	IBM 4381 Model P-21	IBM 4381 Model P-14
CPUs	1	2
Memory capacity	16 million characters	32 million characters
Total disk capacity	180 billion characters, using IBM and Storage Technology disks	
Total disk storage cost	$2 million ($11 per 1 million characters)	
Total cost of IBM mainframes	$1 million	

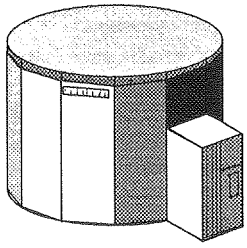

Storage Technology Corp. 4400 Automated Cartridge Systems (ACS)

No. of silos	6
Data-storage capacity	2.1 trillion characters per silo (stored on magnetic tape cartridges)
No. of tape cartridges	6,000 per silo
Cartridge capacity	348 million characters per cartridge (after automatic data compression)
Access time	15 seconds or less to mount a cartridge, 30 seconds (average) to find data
Cartridge read/write drives	22 total (2–4 drives per silo)
Data storage cost	$0.11 per 1 million characters

Mid-Level Computing Systems

Mid-level computers fill the gap between the very powerful, expensive supercomputers and the desktop machines in users' offices. These computers are best suited for highly interactive applications, such as pre- and post-processing of supercomputer data, visualization, software development, and databases.

VAX 8650

Memory capacity	32 million characters
Disk storage capacity	8.5 billion characters (shared with the VAX 6000-320)
Use	Database, mathematical, general computing resource
Manufacturer	Digital Equipment Corp. (DEC)
Machine cost	$400,000

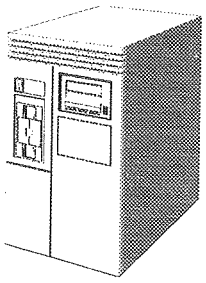

VAX 6000-320

Dual processor Symmetric Multi-Processing machine

Memory capacity	64 million characters
Disk storage capacity	8.5 billion characters (shared with the VAX 8650)
Use	Database, mathematical, general computing resource
Manufacturer	Digital Equipment Corp. (DEC)
Machine cost	$100,000

HP 9000/750

CPU rate	76 MIPS (million instructions per second)
Memory capacity	192 million characters
Disk capacity	9.1 billion characters
CD ROM drive	
DAT (digital-audio tape) drive	1.3 billion characters
Use	Feasibility testing of RISC (reduced instruction set computer) architecture, scalar server, and distributed computing
Manufacturer	Hewlett-Packard Corp.
Machine cost	$95,000

Communications Systems

The Energy Sciences Network (ESnet)

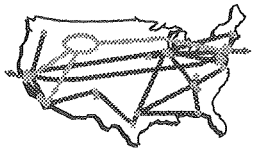

ESnet's role	Funded by DOE/Off. of Energy Research to provide network access to the five major energy programs: Basic Energy Sciences, Health & Environmental Research, High Energy and Nuclear Physics, Magnetic Fusion, and the Superconducting Super Collider
Brought online	Late 1989; replaced MFEnet (Magnetic Fusion Energy Network that had provided access to the Center and its users since 1976)
Sites	
U.S.	32, connected via a T1 (1.536 Mbps) backbone
International	Two sites in Japan (64 and 192 Kbps), one to Brazil (64 Kbps), and one to Germany (256 Kbps); one to Italy planned
Backbone	36 T1 links connecting over 32 cisco multi-protocol routers
Communication protocols	Carried simultaneously: DOD-IP (Dept. of Defense Internet Protocol), DECnet, OSI-CLNP (Open Standards Interconnection Protocol) and X.25
Federal Interagency eXchange	Two FIX locations for data exchange with MILnet, NASA Science Network, NSFnet, and TWB
Regional interconnections	Interconnects with eight regional networks

NERSC Network

Role	High-speed communications connecting all the computers at NERSC
Ethernet (Local Area Network)	10 Mbps
Hyperchannel (built by Network Systems Corp.)	50 Mbps
FDDI (Fiber Distributed Data Interface) High Speed Network	100 Mbps

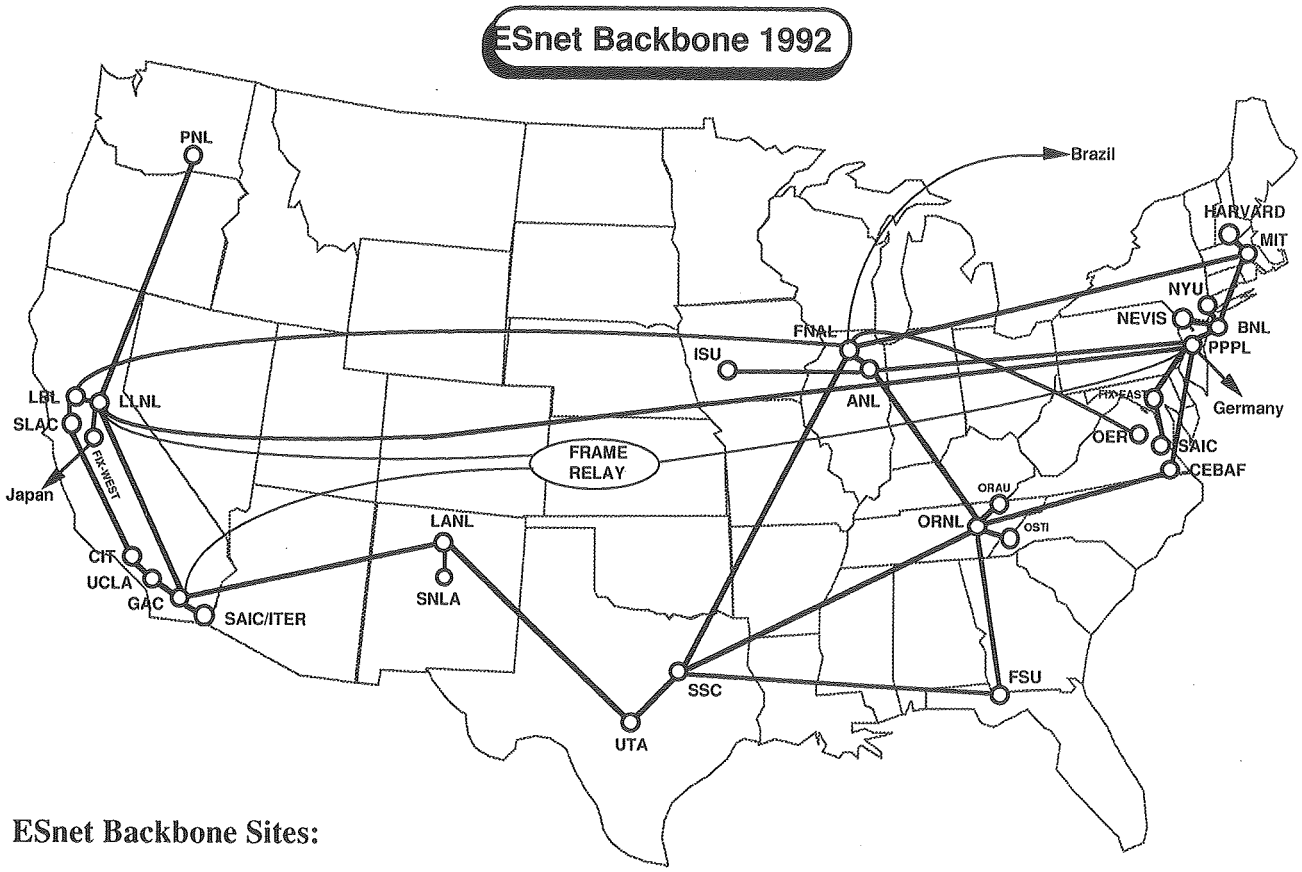

ESnet Backbone Sites:

ANL—Argonne National Lab.
BNL—Brookhaven National Lab.
CEBAF—Continuous Electron Beam Accelerator Facility
CIT—California Institute of Technology
FIX—Federal Interagency eXchange
FNAL—Fermi National Accelerator Lab.
FSU—Florida State Univ. (Supercomputer Comp. Res. Inst.)
GAC—General Atomics Corp.
ISU—Iowa State University (Ames Lab.)
LANL—Los Alamos National Lab.
LBL—Lawrence Berkeley Lab.
LLNL—Lawrence Livermore National Lab.
MIT—Massachusetts Institute of Tech.
NYU—New York University
OER—Office of Energy Research, Wash., D.C.
ORAU—Oak Ridge Associated Universities
ORNL—Oak Ridge National Lab.
OSTI—Office of Scientific and Tech. Info.
PNL—Pacific Northwest Lab.
PPPL—Princeton Plasma Physics Lab.
SAIC—Science Applications, Inc.
SLAC—Stanford Linear Accelerator Lab.
SNLA—Sandia National Lab., Albuquerque
SSC—Superconducting Super Collider
UCLA—University of California, Los Angeles
UTA—Univ. of Texas, Austin (Fusion Res. Center)

ESnet Operations Control Center

Role	**Provides 24 hour/day monitoring and control of various network components that comprise ESnet**
Operated by	**NERSC Engineering Group, Network Special Projects Group, and the Operations staff**
Data collection	**ESnet performance data is collected and saved every 15 minutes for troubleshooting and long-term planning**

Support Services

Machine-Room Operations

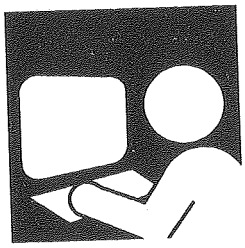

Operator Console Area—Provides the communications and monitoring facilities used by computer operators to keep an eye on the CRAYs, IBMs, shared disk system, the ACSs, and the ESnet communication system. A significant part of the operator's job is diagnosing whether mainframe problems are of hardware or software origin, overseeing the archival data system, and monitoring the ESnet backbone system. *The Machine Room is staffed by operators 24 hours a day, every day of the year.*

Fire Protection—Smoke detectors are located in the ceiling and under the floor. Overhead sprinklers are tripped automatically. Under-the-floor fire-suppressant gas can be discharged by an operator. When smoke is detected, steel overhead fire doors immediately drop behind the lobby windows, and the Laboratory's Fire Department is automatically notified.

Monitoring Panel—The panel in the hall leading to the Machine Room displays information about the environmental and safety conditions in the Machine Room. An alarm sounds if anything is wrong, such as high humidity, cooling system failure, or smoke.

Prime Power, Cooling, and Heating

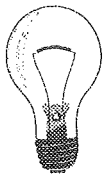

Machine-Room Power Consumption

March 1992	**1,467,446 kilowatt hours**

Cooling

If cooled at a very high rate, computers can have their vital parts very close together and do very fast computations.

CRAY cooling requirements	60 tons (air equivalent) direct cooling per CRAY
CRAY-2s	Directly fluorinert cooled
CRAY X-MP	Refrigerant cooling coils mounted directly in computer assemblies' structural framework
CRAY Y-MP	
Memory and CPU	Cooled by fluorinert, which in turn is cooled by freon
I/O and SSD sections	Cooled directly by freon
Machine Room and building	Two 150-ton air-conditioning systems

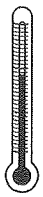

Heating

Building 451 comfort heating	Provided by heat generated in computer room (primarily from disks)

Computer Talk

FLOPs—Supercomputers are measured by how many FLOPs (FLoating point OPerations) they can do in a second. This is typically in the millions (MegaFLOP), but a CRAY-2 can perform over a billion (GigaFLOP), and the CRAY Y-MP and Y-MP C90 (16 GigaFLOPs) can perform even more.

bit—an off–on switch, either 0 or 1; the smallest unit of information in a computer

byte—8 bits or 1 **character**

word—8 bytes or 8 characters on a CRAY; computer memory is divided into words

1 megabyte = 1 million bytes
1 gigabyte = 1 billion bytes
1 terabyte = 1 trillion bytes

Cycle time—Every computer uses a clock to control its rhythm. The speed of this clock's vibrations controls the speed of the machine. CRAY clocks run in the range of a few nanoseconds, where a nanosecond is a billionth of a second. In 1 nanosecond electricity can travel across approximately 1 foot of wire. Each new model gets a faster clock, making the computer perform that much faster.

Believe It or Not

A code taking 96 hours to run on a personal computer would take only 1 minute to run on a supercomputer.

If you did an arithmetic operation like 1 + 1 = 2 at the rate of 60 per minute for every minute in every hour of every day for the next 33 years, you would have accomplished what the CRAY-2 can do in 1 second.

You could do those same arithmetic operations for 535 years and you would have done what the CRAY Y-MP C90 does in 1 second.

ESnet speed—In 1 second it could transfer five 1000-page textbooks in electronic form.

UCRL-TB-111503
September 1992

Work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.